



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/750,601	12/30/2003	Leslie B. Lamport	224009	1522
41505	7590	05/13/2008	EXAMINER	
WOODCOCK WASHBURN LLP (MICROSOFT CORPORATION)			SCIACCA, SCOTT M	
CIRA CENTRE, 12TH FLOOR			ART UNIT	PAPER NUMBER
2929 ARCH STREET			2146	
PHILADELPHIA, PA 19104-2891				

MAIL DATE	DELIVERY MODE
05/13/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/750,601	LAMPORT, LESLIE B.	
	Examiner	Art Unit	
	Scott M. Sciacca	2146	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) Responsive to communication(s) filed on 30 December 2003.
- 2a) This action is FINAL. 2b) This action is non-final.
- 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) Claim(s) 1-39 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) Claim(s) _____ is/are allowed.
- 6) Claim(s) 1-39 is/are rejected.
- 7) Claim(s) _____ is/are objected to.
- 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) The specification is objected to by the Examiner.
- 10) The drawing(s) filed on 30 December 2003 is/are: a) accepted or b) objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) All b) Some * c) None of:
1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ . |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| Paper No(s)/Mail Date <u>12/30/2003, 12/20/2004, 3/7/2005</u> . | 6) <input type="checkbox"/> Other: _____ . |

DETAILED ACTION

Claim Rejections - 35 USC § 101

1. 35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

2. Claims 1-39 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter.

Claims 1-8 are drawn to a method which incorporates a distributed computing algorithm. Claims 9-18 are drawn to a computer-readable medium storing instructions comprising the algorithm. Claims 19-30 are drawn to a computing device capable of executing the algorithm. Claims 31-39 are drawn to the algorithm itself.

Claims 1-39 encompass the abstract idea of voting for values and selecting a value. According to the MPEP, “One may not patent a process that comprises every ‘substantial practical application’ of an abstract idea, because such a patent ‘in practical effect would be a patent on the [abstract idea] itself.’” The MPEP further states “Thus, a claim that recites a computer that solely calculates a mathematical formula (see Benson) or a computer disk that solely stores a mathematical formula is not directed to the type of subject matter eligible for patent protection.” (See MPEP 2106)

Additionally, Claims 9-18 are drawn to a “computer-readable medium”. In accordance with the Applicant’s specification, the computer-readable medium may be

an electromagnetic signal. This subject matter is not limited to that which falls within a statutory category of invention because it is not limited to a process, machine, manufacture, or a composition of matter.

Furthermore, Claims 31-39 are drawn to a “conflict tolerant message delay reducing consensus algorithm”. An algorithm by itself does not fall into a statutory category since it is not limited to a process, machine, manufacture, or a composition of matter.

Claim Rejections - 35 USC § 102

3. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

4. Claims 31-37 and 39 are rejected under 35 U.S.C. 102(b) as being anticipated by US 5,261,085 (hereinafter, ‘085).

Regarding Claim 31, ‘085 teaches a conflict tolerant message delay reducing consensus algorithm for use in computing environment comprising a dedicated client device and a distributed computing system implemented by one or more devices, wherein the computing system implements the conflict tolerant message delay reducing

consensus algorithm, the conflict tolerant message delay reducing consensus algorithm comprising:

transmitting one or more proposed values from one or more clients (“*One of the processes in the network is designated as a leader, and it sends ballots with proposed commands to the other processes*” – See Col. 3, lines 14-16);

voting, at one or more of the one or more devices implementing the distributed computing system, for a proposed value from among the one or more proposed values (“*In each ballot, a process has the choice of either voting for the proposed command or not voting*” – See Col. 3, lines 16-18), wherein the proposed value was proposed by a client having a most dominant client identifier from among the one or more clients proposing values (“*One suitable method of choosing the leader is to select the process with the highest identification number*” – See Col. 8, lines 33-34);

transmitting to one or more of the one or more devices implementing the distributed computing system an indication of the vote for the proposed value (“*MaxVote(b, q, β)_{dec} is the vote with the largest ballot number less than b among all the votes cast by process q , or null _{q} if process q did not vote in any ballot numbered less than b . The leader obtains MaxVote(b, q, β)_{dec} from each process q by an exchange of messages*” – See Col. 4, lines 51-55); and

transmitting, to the client having the highest client identifier, a result of the vote for the proposed value (“*The preliminary protocol for conducting a single ballot initiated by process p is follows*” – See Col. 4, lines 56-57; “*A process q responds to the receipt*

of the NextBallot(b) message by sending a LastVote(b,v) message to p” – See Col. 4, lines 61-62).

Regarding Claim 32, ‘085 further teaches the one or more devices implementing the distributed computing system also acting as clients of the distributed computing system (“*a distributed system 11 in which the invention is employed has a set of computers 12 which are connected together by a network 13*” – See Col. 3, lines 1-3).

Regarding Claim 33, ‘085 further teaches the dedicated client device being identified by a least dominant client identifier (“*The selection requirement is then satisfied by having one of the processes send a message containing its identification number to every other process*” – See Col. 8, lines 35-37).

Regarding Claim 34, ‘085 further teaches determining that the distributed computing system has selected the proposed value when each of the one or more devices implementing the distributed computing system has voted for the proposed value (“*In order for a ballot to succeed and a command to be issued, a majority set of the processes in the system must vote for it*” – See Col. 3, lines 18-20).

Regarding Claim 35, ‘085 further teaches the proposed value comprising a function, and wherein the voting for the proposed value comprises provisionally executing the function in a system step (“*To be issued, a command must be voted for*

by a majority of the processes in the system. Each issued command is stored by each of the processes in the majority set which voted for it” – See Col. 2, lines 37-40).

Regarding Claim 36, ‘085 teaches the voting for the proposed value comprising changing a previous vote if the previous vote was for a previously proposed value, proposed by a previous client having a client identifier that is less dominant than the client proposing the proposed value (“*the receipt of a NextBallot(b) message by a process q in step 2 may set nextbal(q) to a value that prevents q from voting in step 4 for any previously initiated ballot*” – See Col. 7, lines 31-33).

Regarding Claim 37, ‘085 teaches ending the conflict tolerant message delay reducing consensus algorithm and commencing a fault tolerant consensus algorithm if a failure is detected (“*The invention provides a system and method for implementing a distributed state machine in which consistency is maintained despite the failure of any number of processes and communication paths*” – See Col. 2, lines 22-25).

Regarding Claim 39, ‘085 further teaches identifying a possibly selected proposed value, wherein the possibly selected proposed value is any value if at least one of the one or more devices implementing the distributed computing system did not previously vote (“*any command for which one of the processes has previously voted but does not have a command number is broadcast as a proposed command in a new ballot*” – See Col. 2, lines 48-51).

Claim Rejections - 35 USC § 103

5. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

6. Claims 1-4, 6-14, 16-24 and 26-30 are rejected under 35 U.S.C. 103(a) as being unpatentable over US 5,261,085 (hereinafter, ‘085) in view of Paxos Made Simple (hereinafter, Paxos), from Applicant’s IDS dated December 30, 2003.

Regarding Claim 1, ‘085 teaches a method for selecting a value in a distributed computing system, the method comprising:

receiving a first proposed value from a first client (“*One of the processes in the network is designated as a leader, and it sends ballots with proposed commands to the other processes*” – See Col. 3, lines 14-16) having a first client identifier (“*The*

preliminary protocol for conducting a single ballot initiated by process p is follows” – See Col. 4, lines 56-57);

voting for the first proposed value in a first system step (“*In each ballot, a process has the choice of either voting for the proposed command or not voting*” – See Col. 3, lines 16-18);

transmitting a first indication of the voting for the first proposed value to one or more devices (“*MaxVote(b,q, β)_{dec} is the vote with the largest ballot number less than b among all the votes cast by process q, or null_q if process q did not vote in any ballot numbered less than b. The leader obtains MaxVote(b,q, β)_{dec} from each process q by an exchange of messages*” – See Col. 4, lines 51-55); and

transmitting a first result of the voting for the first proposed value to the first client (“*A process q responds to the receipt of the NextBallot(b) message by sending a LastVote(b,v) message to p*” – See Col. 4, lines 61-62).

‘085 teaches clients having identifiers and determining if a second identifier is more dominant than a first identifier (“*One suitable method of choosing the leader is to select the process with the highest identification number*” – See Col. 8, lines 33-34).

‘085 does not explicitly teach that the voting for the first proposed value, the transmitting the first indication of the voting for the first proposed value, and the transmitting the first result are not performed if a second proposed value was proposed by a second client having a second client identifier that is more dominant than the first client identifier and the second proposed value was previously voted for.

However, Paxos discloses that a proposal numbered n is “accepted” by an acceptor if and only if the acceptor has not already accepted a request having a number greater (more dominant) than n (“*A proposer sends a proposed value to a set of acceptors. An acceptor may accept the proposed value. The value is chosen when a large enough set of acceptors have accepted it*” – See p. 2, lines 13-15; “*An acceptor can accept a proposal numbered n iff it has not responded to a prepare request having a number greater than n* ” – See p. 5, lines 10-11). Paxos describes an acceptor accepting a proposed value in a fashion similar to the way ‘085 describes voting for a proposed value. It would have been obvious to one of ordinary skill in the art at the time the invention was made to modify the method of selecting a value disclosed by ‘085 to include not voting for a first proposed value, not transmitting a first indication of the voting for the first proposed value, and not transmitting a first result if a second proposed value was proposed by a second client having a second client identifier that is more dominant than the first client identifier and the second proposed value was previously voted for. One of ordinary skill could easily apply the principle of comparing proposal numbers shown by Paxos to the teachings of ‘085 to end up with comparing identifiers of a first and second client, each of whom generates a proposal. Paxos describes a number of “safety requirements” for guaranteeing consensus in a distributed system on p. 1, lines 18-23. The step of comparing proposal numbers along with the respective outcomes of the comparison is necessary in order to satisfy the safety properties which guarantee consensus (See p. 5, lines 13-14).

Regarding Claim 2, '085 further teaches that the first proposed value comprises a first function, and wherein the voting for the first proposed value comprises provisionally executing the first function in the first system step ("To be issued, a command must be voted for by a majority of the processes in the system. Each issued command is stored by each of the processes in the majority set which voted for it" – See Col. 2, lines 37-40).

Regarding Claim 3, '085 further teaches the voting for the first proposed value comprising changing a previous vote for the second proposed value if the second proposed value was previously voted for and if the second client identifier is less dominant than the first client identifier ("the receipt of a NextBallot(b) message by a process q in step 2 may set $nextbal(q)$ to a value that prevents q from voting in step 4 for any previously initiated ballot" – See Col. 7, lines 31-33).

Regarding Claim 4, '085 further teaches the first proposed value comprising a first function identified by a first function identifier ("one of the processes send a message containing its identification number to every other process" – See Col. 8, lines 36-37), and wherein the voting for the first proposed value comprises executing the first function in the first system step ("Steps 1-4 thus contain the protocol for initiating a ballot and voting on it. In step 5, the results of the balloting are determined, and in step 6 the command is declared to be issued" – See Col. 5, lines 59-62) unless the first function identifier is equivalent to a second function identifier that identifies a second

function, wherein the second function was executed in a second system step that preceded the first system step (“*Receiving multiple copies of a message can cause an action to be repeated. Except in step 3, performing the action a second time has no effect. For example, sending several Voted(b,q) messages in step 4 has the same effect as sending just one. The repetition of step 3 is prevented by using the entry made in stable storage when it is executed. Thus, the consistency condition is maintained even if the same message is received several times*” – See Col. 5, lines 50-58).

Regarding Claim 6, ‘085 further teaches:

receiving a message (“*The network can be of any suitable type or configuration which permits messages to be sent between any two computers on the network*” – See Col. 3, lines 9-11), wherein the message is part of a fault tolerant consensus algorithm (“*This invention pertains generally to distributed computing systems and, more particularly, to a distributed system and method utilizing a state machine approach and having the ability to continue working despite the occurrence of a fault such as a processor stopping or running slowly*” – See Col. 1, lines 8-13);

ignoring additional proposed values from the first client (“*For example, a process q in the majority set can ignore a NextBallot(b) message*” – See Col. 5, lines 41-42); and participating in the fault tolerant consensus algorithm (“*consistency is maintained despite the failure of any number of processes and communication paths*” – See Col. 10, lines 64-66).

Regarding Claim 7, ‘085 further teaches the participating in the fault tolerant consensus algorithm comprising transmitting a possibly selected proposed value if a proposed value was previously voted for, wherein the possibly selected proposed value was previously voted for and was proposed by a client having a most dominant client identifier among all clients whose proposals were received and who proposed values for a current system step (“*As part of this procedure, any command for which one of the processes has previously voted but does not have a command number is broadcast as a proposed command in a new ballot*” – See Col. 2, lines 48-51).

Regarding Claim 8, ‘085 further teaches:

transmitting one or more polling messages to initiate a fault tolerant consensus algorithm (“*The preliminary protocol for conducting a single ballot initiated by process p is follows: 1. Process p (the leader) chooses a new ballot number b and sends a NextBallot(b) message to some set of processes*” – See Col. 4, lines 56-60);

receiving one or more vote indication messages in response to the one or more polling messages (“*2. A process q responds to the receipt of the NextBallot(b) message by sending a LastVote(b,v) message to p*” – See Col. 4, lines 61-62); and

selecting, as a third proposed value, any value if the one or more vote indication messages indicate that at least one device has not previously voted or if the one or more vote indication messages indicate two or more different possibly selected proposed values (“*any command for which one of the processes has previously voted but does not have a command number is broadcast as a proposed command in a new ballot*” – See Col. 4, lines 63-67).

ballot" – See Col. 2, lines 48-51), wherein a possibly selected proposed value was previously voted for by a device and was proposed by a client having a most dominant client identifier among all clients whose proposals were received by the device, and wherein further the third proposed value is proposed using the fault tolerant consensus algorithm ("The selection requirement is then satisfied by having one of the processes send a message containing its identification number to every other process every T-11 msec, for some suitable choice of T" – See Col. 8, lines 35-38).

Regarding Claim 9, '085 teaches a computer-readable medium having computer-executable instructions for selecting a value in a distributed computing system, the computer-executable instructions performing steps comprising:

receiving a first proposed value ("One of the processes in the network is designated as a leader, and it sends ballots with proposed commands to the other processes" – See Col. 3, lines 14-16) from a first client having a first client identifier ("The preliminary protocol for conducting a single ballot initiated by process p is follows" – See Col. 4, lines 56-57);

voting for the first proposed value in a first system step ("In each ballot, a process has the choice of either voting for the proposed command or not voting" – See Col. 3, lines 16-18);

transmitting a first indication of the voting for the first proposed value to one or more devices ("MaxVote(b, q, β)_{dec} is the vote with the largest ballot number less than b among all the votes cast by process q, or null_q if process q did not vote in any ballot

numbered less than b. The leader obtains $\text{MaxVote}(b, q, \beta)_{\text{dec}}$ from each process q by an exchange of messages” – See Col. 4, lines 51-55); and

transmitting a first result of the voting for the first proposed value to the first client (“*A process q responds to the receipt of the $\text{NextBallot}(b)$ message by sending a $\text{LastVote}(b, v)$ message to p*” – See Col. 4, lines 61-62).

‘085 teaches clients having identifiers and determining if a second identifier is more dominant than a first identifier (“*One suitable method of choosing the leader is to select the process with the highest identification number*” – See Col. 8, lines 33-34). ‘085 does not explicitly teach that the voting for the first proposed value, the transmitting the first indication of the voting for the first proposed value, and the transmitting the first result are not performed if a second proposed value was proposed by a second client having a second client identifier that is more dominant than the first client identifier and the second proposed value was previously voted for.

However, Paxos discloses that a proposal numbered n is “accepted” by an acceptor if and only if the acceptor has not already accepted a request having a number greater (more dominant) than n (“*A proposer sends a proposed value to a set of acceptors. An acceptor may accept the proposed value. The value is chosen when a large enough set of acceptors have accepted it*” – See p. 2, lines 13-15; “*An acceptor can accept a proposal numbered n iff it has not responded to a prepare request having a number greater than n*” – See p. 5, lines 10-11). Paxos describes an acceptor accepting a proposed value in a fashion similar to the way ‘085 describes voting for a proposed value. It would have been obvious to one of ordinary skill in the art at the time

the invention was made to modify the steps of selecting a value disclosed by '085 to include not voting for a first proposed value, not transmitting a first indication of the voting for the first proposed value, and not transmitting a first result if a second proposed value was proposed by a second client having a second client identifier that is more dominant than the first client identifier and the second proposed value was previously voted for the same reasons as those given with respect to Claim 1.

Regarding Claim 10, '085 further teaches that the first proposed value comprises a first function, and wherein the voting for the first proposed value comprises provisionally executing the first function in the first system step ("To be issued, a command must be voted for by a majority of the processes in the system. Each issued command is stored by each of the processes in the majority set which voted for it" – See Col. 2, lines 37-40).

Regarding Claim 11, '085 further teaches the voting for the first proposed value comprising changing a previous vote for the second proposed value if the second proposed value was previously voted for and if the second client identifier is less dominant than the first client identifier ("the receipt of a NextBallot(b) message by a process q in step 2 may set nextbal(q) to a value that prevents q from voting in step 4 for any previously initiated ballot" – See Col. 7, lines 31-33).

Regarding Claim 12, '085 further teaches the second proposed value comprising a second proposed function, and wherein the changing the previous vote comprises undoing a previous execution of the second proposed function ("a process has the option not to vote, even if casting a vote is not precluded by a previous LastVote message" – See Col. 5, lines 38-40).

Regarding Claim 13, '085 further teaches the second proposed value comprises a second proposed function, and wherein the changing the previous vote comprises allowing a previous provisional execution of the second proposed function to expire ("the initiation of a new ballot can prevent any previously initiated ballot from succeeding" – See Col. 7, lines 34-35).

Regarding Claim 14, '085 further teaches the first proposed value comprising a first function identified by a first function identifier ("one of the processes send a message containing its identification number to every other process" – See Col. 8, lines 36-37), and wherein the voting for the first proposed value comprises executing the first function in the first system step ("Steps 1-4 thus contain the protocol for initiating a ballot and voting on it. In step 5, the results of the balloting are determined, and in step 6 the command is declared to be issued" – See Col. 5, lines 59-62) unless the first function identifier is equivalent to a second function identifier that identifies a second function, wherein the second function was executed in a second system step that preceded the first system step ("Receiving multiple copies of a message can cause an

action to be repeated. Except in step 3, performing the action a second time has no effect. For example, sending several Voted(b,q) messages in step 4 has the same effect as sending just one. The repetition of step 3 is prevented by using the entry made in stable storage when it is executed. Thus, the consistency condition is maintained even if the same message is received several times” – See Col. 5, lines 50-58).

Regarding Claim 16, '085 further teaches:

receiving a message (“The network can be of any suitable type or configuration which permits messages to be sent between any two computers on the network” – See Col. 3, lines 9-11), wherein the message is part of a fault tolerant consensus algorithm (“This invention pertains generally to distributed computing systems and, more particularly, to a distributed system and method utilizing a state machine approach and having the ability to continue working despite the occurrence of a fault such as a processor stopping or running slowly” – See Col. 1, lines 8-13);

ignoring additional proposed values from the first client (“For example, a process q in the majority set can ignore a NextBallot(b) message” – See Col. 5, lines 41-42); and participating in the fault tolerant consensus algorithm (“consistency is maintained despite the failure of any number of processes and communication paths” – See Col. 10, lines 64-66).

Regarding Claim 17, '085 further teaches the participating in the fault tolerant consensus algorithm comprising transmitting a possibly selected proposed value if a

proposed value was previously voted for, wherein the possibly selected proposed value was previously voted for and was proposed by a client having a most dominant client identifier among all clients whose proposals were received and who proposed values for a current system step (“*As part of this procedure, any command for which one of the processes has previously voted but does not have a command number is broadcast as a proposed command in a new ballot*” – See Col. 2, lines 48-51).

Regarding Claim 18, ‘085 further teaches:

transmitting one or more polling messages to initiate a fault tolerant consensus algorithm (“*The preliminary protocol for conducting a single ballot initiated by process p is follows: 1. Process p (the leader) chooses a new ballot number b and sends a NextBallot(b) message to some set of processes*” – See Col. 4, lines 56-60);

receiving one or more vote indication messages in response to the one or more polling messages (“*2. A process q responds to the receipt of the NextBallot(b) message by sending a LastVote(b,v) message to p*” – See Col. 4, lines 61-62); and

selecting, as a third proposed value, any value if the one or more vote indication messages indicate that at least one device has not previously voted or if the one or more vote indication messages indicate two or more different possibly selected proposed values (“*any command for which one of the processes has previously voted but does not have a command number is broadcast as a proposed command in a new ballot*” – See Col. 2, lines 48-51), wherein a possibly selected proposed value was previously voted for by a device and was proposed by a client having a most dominant

client identifier among all clients whose proposals were received by the device, and wherein further the third proposed value is proposed using the fault tolerant consensus algorithm (“*The selection requirement is then satisfied by having one of the processes send a message containing its identification number to every other process every T-11 msec, for some suitable choice of T*” – See Col. 8, lines 35-38).

Regarding Claim 19, ‘085 teaches a computing device operating as part of a distributed computing system (Computer 12 – See Fig. 1), the computing device comprising a processing unit (“*Each of the computers includes at least a processor*” – See Col. 3, lines 3-4) and a network interface (“*a distributed system 11 in which the invention is employed has a set of computers 12 which are connected together by a network 13*” – See Col. 3, lines 1-3).

‘085 teaches the network interface performing the steps comprising: receiving the first proposed value from a first client (“*One of the processes in the network is designated as a leader, and it sends ballots with proposed commands to the other processes*” – See Col. 3, lines 14-16) having the first client identifier (“*The preliminary protocol for conducting a single ballot initiated by process p is follows*” – See Col. 4, lines 56-57);

transmitting a first indication of the voting for the first proposed value to one or more devices also operating as part of the distributed computing system (“*MaxVote(b,q, β)_{dec} is the vote with the largest ballot number less than b among all the votes cast by process q, or null_q if process q did not vote in any ballot numbered less*

than b. The leader obtains $\text{MaxVote}(b, q, \beta)_{\text{dec}}$ from each process q by an exchange of messages” – See Col. 4, lines 51-55); and

transmitting a first result of the voting for the first proposed value to the first client (“*A process q responds to the receipt of the $\text{NextBallot}(b)$ message by sending a $\text{LastVote}(b, v)$ message to p*” – See Col. 4, lines 61-62).

‘085 teaches clients having identifiers and comparing identifiers to determine if a second identifier is more dominant than a first identifier (“*One suitable method of choosing the leader is to select the process with the highest identification number*” – See Col. 8, lines 33-34). ‘085 does not explicitly teach that the comparing is performed if a second proposed value, proposed by a second client having the second client identifier, was previously voted for in a first system step. However, Paxos teaches comparing identifiers for proposed values from a first and second client and voting for a first proposed value if the first client identifier is more dominant than the second client identifier (“*A proposer sends a proposed value to a set of acceptors. An acceptor may accept the proposed value. The value is chosen when a large enough set of acceptors have accepted it*” – See p. 2, lines 13-15; “*An acceptor can accept a proposal numbered n iff it has not responded to a prepare request having a number greater than n*” – See p. 5, lines 10-11).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to modify the steps of selecting a value disclosed by ‘085 to include comparing a first client identifier to a second client identifier if a second proposed value, proposed by a second client having the second client identifier, was previously voted for

in a first system step and voting for a first proposed value in the first system step if the first client identifier is more dominant than the second client identifier and the second proposed value was previously voted for the same reasons as those given with respect to Claim 1.

Regarding Claim 20, '085 further teaches the first proposed value comprising a first function, and wherein the voting for the first proposed value comprises provisionally executing the first function in the first system step ("To be issued, a command must be voted for by a majority of the processes in the system. Each issued command is stored by each of the processes in the majority set which voted for it" – See Col. 2, lines 37-40).

Regarding Claim 21, '085 further teaches the voting for the first proposed value comprising changing a previous vote for the second proposed value if the second proposed value was previously voted for and if the second client identifier is less dominant than the first client identifier ("the receipt of a NextBallot(b) message by a process q in step 2 may set $nextbal(q)$ to a value that prevents q from voting in step 4 for any previously initiated ballot" – See Col. 7, lines 31-33).

Regarding Claim 22, '085 teaches the second proposed value comprising a second proposed function, and wherein the changing the previous vote comprises undoing a previous execution of the second proposed function ("a process has the

option not to vote, even if casting a vote is not precluded by a previous LastVote message” – See Col. 5, lines 38-40).

Regarding Claim 23, ‘085 further teaches the second proposed value comprising a second proposed function, and wherein the changing the previous vote comprises allowing a previous provisional execution of the second proposed function to expire (“*the initiation of a new ballot can prevent any previously initiated ballot from succeeding*” – See Col. 7, lines 34-35).

Regarding Claim 24, ‘085 teaches the first proposed value comprising a first function identified by a first function identifier (“*one of the processes send a message containing its identification number to every other process*” – See Col. 8, lines 36-37), and wherein the voting for the first proposed value comprises executing the first function in the first system step (“*Steps 1-4 thus contain the protocol for initiating a ballot and voting on it. In step 5, the results of the balloting are determined, and in step 6 the command is declared to be issued*” – See Col. 5, lines 59-62) unless the first function identifier is equivalent to a second function identifier that identifies a second function, wherein the second function was executed in a second system step that preceded the first system step (“*Receiving multiple copies of a message can cause an action to be repeated. Except in step 3, performing the action a second time has no effect. For example, sending several Voted(b,q) messages in step 4 has the same effect as sending just one. The repetition of step 3 is prevented by using the entry made in stable*”

storage when it is executed. Thus, the consistency condition is maintained even if the same message is received several times” – See Col. 5, lines 50-58).

Regarding Claim 26, ‘085 further teaches:

ignoring additional proposed values from the first client (“*For example, a process q in the majority set can ignore a NextBallot(b) message*” – See Col. 5, lines 41-42); and participating in a fault tolerant consensus algorithm (“*consistency is maintained despite the failure of any number of processes and communication paths*” – See Col. 10, lines 64-66); and

wherein the network interface performs further steps comprising:
receiving a message (“*The network can be of any suitable type or configuration which permits messages to be sent between any two computers on the network*” – See Col. 3, lines 9-11), wherein the message is part of the fault tolerant consensus algorithm (“*This invention pertains generally to distributed computing systems and, more particularly, to a distributed system and method utilizing a state machine approach and having the ability to continue working despite the occurrence of a fault such as a processor stopping or running slowly*” – See Col. 1, lines 8-13).

Regarding Claim 27, ‘085 further teaches the participating in the fault tolerant consensus algorithm comprising transmitting a possibly selected proposed value if a proposed value was previously voted for, wherein the possibly selected proposed value was previously voted for and was proposed by a client having a most dominant client

identifier among all clients who proposed values to the computing device for a current system step (“*As part of this procedure, any command for which one of the processes has previously voted but does not have a command number is broadcast as a proposed command in a new ballot*” – See Col. 2, lines 48-51).

Regarding Claim 28, ‘085 further teaches:

selecting, as a third proposed value, any value if one or more vote indication messages indicate that at least one device has not previously voted or if the one or more vote indication messages indicate two or more different possibly selected proposed values (“*any command for which one of the processes has previously voted but does not have a command number is broadcast as a proposed command in a new ballot*” – See Col. 2, lines 48-51),

wherein a possibly selected proposed value was previously voted for by a device and was proposed by a client having a most dominant client identifier among all clients whose proposals were received by the device, and wherein further the third proposed value is proposed using the fault tolerant consensus algorithm (“*The selection requirement is then satisfied by having one of the processes send a message containing its identification number to every other process every T-11 msec, for some suitable choice of T*” – See Col. 8, lines 35-38); and

wherein the network interface performs further steps comprising:
transmitting one or more polling messages to initiate the fault tolerant consensus algorithm (“*The preliminary protocol for conducting a single ballot initiated by process p*

is follows: 1. Process p (the leader) chooses a new ballot number b and sends a NextBallot(b) message to some set of processes” – See Col. 4, lines 56-60); and receiving the one or more vote indication messages in response to the one or more polling messages (“2. A process q responds to the receipt of the NextBallot(b) message by sending a LastVote(b,v) message to p” – See Col. 4, lines 61-62).

Regarding Claim 29, '085 further teaches the operating as part of the distributed computing system comprising operating as a client of the distributed computing system (“*a distributed system 11 in which the invention is employed has a set of computers 12 which are connected together by a network 13*” – See Col. 3, lines 1-3).

Regarding Claim 30, '085 further teaches the distributed computing system being comprised of devices that are also clients of the distributed computing system (“*a distributed system 11 in which the invention is employed has a set of computers 12 which are connected together by a network 13*” – See Col. 3, lines 1-3).

7. Claims 5, 15 and 25 are rejected under 35 U.S.C. 103(a) as being unpatentable over US 5,261,085 (hereinafter, '085) in view of Paxos Made Simple (hereinafter, Paxos) and further in view of US 2003/0227392 (hereinafter, '392).

Regarding Claim 5, '085 in view of Paxos teaches the method of Claim 1. '085 and Paxos do not explicitly teach the first proposed value comprising a first idempotent

function, and wherein the voting for the first proposed value comprises executing the first idempotent function in the first system step even if the first idempotent function is equivalent to a second idempotent function that was executed in a second system step that preceded the first system step.

However, '392 discloses the use of idempotent operations ("In general it is not easy to be certain which events were acted upon and which were lost. Furthermore, there is no inherent ordering of the processing of events through the real time semantics section 1630 and the core 1610. Replaying sufficient events to cover the largest possible loss works, provided all processing is idempotent and order independent" – See [0263]) so that a first function idempotent function may be executed even though an equivalent idempotent function has been executed previously ("If this condition is met, the semantics of the operations invoked 2060 are such that if events are processed in different orders or if the same event is processed more than once, the ultimate system state will be identical" – See [0263]).

It would have been obvious to one of ordinary skill at the time the invention was made to use idempotent functions with the distributed computing method taught by '085 and Paxos. '085 and Paxos generally disclose distributed computing systems where commands (functions) are distributed for execution by different machines on a network. '392 teaches that when idempotent functions are used, if the same function is processed (executed) more than once, then the end result will be identical (See [0263]).

Regarding Claim 15, '085 in view of Paxos teaches the computer-readable medium of Claim 9. '085 and Paxos do not explicitly teach the first proposed value comprising a first idempotent function, and wherein the voting for the first proposed value comprises executing the first idempotent function in the first system step even if the first idempotent function is equivalent to a second idempotent function that was executed in a second system step that preceded the first system step.

However, '392 discloses the use of idempotent operations (*"In general it is not easy to be certain which events were acted upon and which were lost. Furthermore, there is no inherent ordering of the processing of events through the real time semantics section 1630 and the core 1610. Replaying sufficient events to cover the largest possible loss works, provided all processing is idempotent and order independent"* – See [0263]) so that a first function idempotent function may be executed even though an equivalent idempotent function has been executed previously (*"If this condition is met, the semantics of the operations invoked 2060 are such that if events are processed in different orders or if the same event is processed more than once, the ultimate system state will be identical"* – See [0263]).

It would have been obvious to one of ordinary skill at the time the invention was made to use idempotent functions with the distributed computing method taught by '085 and Paxos for the same reasons as those given with respect to Claim 5.

Regarding Claim 25, '085 in view of Paxos teaches the computing device of Claim 19. '085 and Paxos do not explicitly teach the first proposed value comprising

a first idempotent function, and wherein the voting for the first proposed value comprises executing the first idempotent function in the first system step even if the first idempotent function is equivalent to a second idempotent function that was executed in a second system step that preceded the first system step.

However, '392 discloses the use of idempotent operations ("In general it is not easy to be certain which events were acted upon and which were lost. Furthermore, there is no inherent ordering of the processing of events through the real time semantics section 1630 and the core 1610. Replaying sufficient events to cover the largest possible loss works, provided all processing is idempotent and order independent" – See [0263]) so that a first function idempotent function may be executed even though an equivalent idempotent function has been executed previously ("If this condition is met, the semantics of the operations invoked 2060 are such that if events are processed in different orders or if the same event is processed more than once, the ultimate system state will be identical" – See [0263]).

It would have been obvious to one of ordinary skill at the time the invention was made to use idempotent functions with the distributed computing method taught by '085 and Paxos for the same reasons as those given with respect to Claim 5.

8. Claim 38 is rejected under 35 U.S.C. 103(a) as being unpatentable over US 5,261,085 (hereinafter, '085) in view of US 2002/0112198 (hereinafter, '198).

Regarding Claim 38, '085 does not explicitly teach the computing environment comprising a monitoring device and the failure being detected by a monitoring device. However, '198 does teach a monitoring device that detects a failure ("At one extreme, the system administrator, operator or a third party (e.g., software for monitoring devices) detects a device failure" – See [0058]). It would have been obvious to one of ordinary skill in the art at the time the invention was made to include a monitoring device within the computing environment taught by '085. Motivation for doing so would be to facilitate recovery in the event that a device fails (See paragraph [0058] of '198).

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Scott M. Sciacca whose telephone number is (571) 270-1919. The examiner can normally be reached on Monday thru Friday, 7:30 A.M. - 5:00 P.M. EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Jeff Pwu can be reached on (571) 272-6798. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/S. M. S./
Examiner, Art Unit 2146
/John Follansbee/
Supervisory Patent Examiner, Art Unit 2151